

Serial No. 09/777,139

Reply to Final Office Action of January 5, 2006

Amendments to the Claims:

This listing of claims will replace all prior versions and listings of claims in the application:

Listing of Claims:

1. (previously presented) A data storage system comprising:
a plurality of storage nodes;
data storage mechanisms implemented in each storage node;
a communication medium linking storage nodes; and
data distributed across a set of the storage nodes selected based on state information corresponding to the storage nodes such that the data remains available irrespective of the unavailability of one or more of the storage nodes within the selected set, wherein the data storage mechanisms on at least two storage nodes collectively implement a unitary volume of network storage.

Claim 2 (canceled)

3. (original) The data storage system of claim 1 wherein the communication medium comprises:
a public network for receiving access requests for the data storage system;
and
a private network enabling communication between storage nodes.

4. (original) The data storage system of claim 3 wherein the public network comprises the Internet.

5. (original) The data storage system of claim 3 wherein the private network comprises a virtual private network implemented over the Internet.

Serial No. 09/777,139

Reply to Final Office Action of January 5, 2006

6. (original) The data storage system of claim 1 further comprising:
communication processes implemented within each of the storage nodes
operable to exchange state information between at least some of the other data
storage nodes.

7. (original) The data storage system of claim 1 wherein each of the
data storage nodes further comprises data structures configured to store state
information about one or more other nodes and the communication links between
them.

8. (original) The data storage system of claim 7 wherein the state
information comprises information selected from the group consisting of but not
limited to: availability information, capacity information, quality of service
information, performance information, geographical location information, network
topological location information.

9. (original) The data storage system of claim 8 wherein the set of
storage nodes is selected by a first of the storage nodes using the state information
stored in the first of the storage nodes.

10. (previously presented) A data storage system comprising:
a plurality of storage nodes wherein each of the data storage nodes further
comprises data structures configured to store state information about one or more
other nodes and the communication links between them;
data storage mechanisms implemented in each storage node;
a communication medium linking storage nodes;
communication processes implemented within each of the storage nodes
operable to exchange state information between at least some of the other data
storage nodes;

Serial No. 09/777,139

Reply to Final Office Action of January 5, 2006

data distributed across a selected set of the storage nodes such that the data remains available irrespective of the unavailability of one or more of the storage nodes within the selected set; and

wherein the communication processes implement a repetitive peer-to-peer conversation between the set of storage nodes enabling the state information contained in the state information data structures within each individual node to represent a consistent view of the state of the collection of storage nodes.

11. (original) The data storage system of claim 1 wherein the network comprises:

a plurality of first level networks, each first level network coupling multiple storage nodes; and

a second level network coupling at least two of the first level networks.

12. (original) The data storage system of claim 11 wherein the first level network comprises a connection selected from the group consisting of: Ethernet, fast Ethernet, gigabit Ethernet, Fibre channel, ATM, firewire, Myrnet, SCSI, serial, parallel, universal serial bus, and wireless networks.

13. (original) The data storage system of claim 1 further comprising:
storage management processes executing on one of the storage nodes to determine state information about each of the set of storage nodes.

14. (original) The data storage system of claim 1 wherein the communication medium comprises a secure communication medium.

15. (original) The data storage system of claim 1 wherein the communication medium implements an authentication protocol between linked storage nodes.

Serial No. 09/777,139

Reply to Final Office Action of January 5, 2006

16. (original) The data storage system of claim 1 wherein the communication medium implements cryptographic security between linked storage nodes.

17. (previously presented) A method of managing data storage in a network comprising multiple storage nodes, the method comprising the acts of:
communicating a storage request to at least one storage node;
causing the at least one storage node to implement the storage request using an arbitrary subset of the storage nodes;
communicating state information between the multiple storage nodes; and
selecting the arbitrary subset of the multiple storage nodes to be used based upon the state information.

Claim 18 (cancelled)

19. (original) The method of claim 17 wherein the act of implementing the storage request comprises associating error checking and correcting (ECC) code with storage request.

20. (original) The method of claim 19 wherein the ECC code is stored in a single network storage node and the unit of data is stored in two or more network storage nodes.

21. (original) The method of claim 17 further comprising:
retrieving a stored unit of data specified by the storage request; and
verifying the correctness of the stored unit of data;
upon detection of an error in the retrieved unit of data, retrieving the correct unit of data using data stored in the others of the arbitrary subset of the multiple storage nodes.

Serial No. 09/777,139

Reply to Final Office Action of January 5, 2006

22. (previously presented) A method of managing data storage in a network comprising multiple storage nodes, the method comprising:

- communicating a storage request to at least one storage node;
- causing the at least one storage node to implement the storage request using an arbitrary subset of the storage nodes;
- attempting to retrieve the stored unit of data from the arbitrary subset of the multiple storage nodes;
- detecting unavailability of one or more network storage nodes; and
- in response to detected unavailability, retrieving the correct unit of data using data stored in others of the arbitrary subset of the multiple storage nodes.

23. (original) The method of claim 22 wherein the unavailability is caused by failure of one or more of the network storage nodes.

24. (original) The method of claim 22 wherein the unavailability is caused by congestion/failure of a network link leading to one or more of the network storage nodes.

25. (original) The method of claim 17 further comprising moving the stored unit of data from one network storage node to another network storage node after the step of storing.

26. (original) The method of claim 17 further comprising:

- communicating state information and storage requests amongst the arbitrary subset of the storage nodes; and
- encrypting at least some of the information and storage requests before communicating them between storage nodes.

27. (original) The method of claim 17 further comprising:

- communicating state information and storage requests amongst the arbitrary subset of the storage nodes; and

Serial No. 09/777,139

Reply to Final Office Action of January 5, 2006

authenticating the communication between storage nodes.

Claims 28-36 (cancelled)

37. (new) A system for managing data storage in a network comprising multiple storage nodes, comprising:

means for communicating a storage request to at least one storage node;

means for causing the at least one storage node to implement the storage request using an arbitrary subset of the storage nodes;

means for communicating state information between the multiple storage nodes; and

means for selecting the arbitrary subset of the multiple storage nodes to be used based upon the state information.

38. (new) The system of claim 37, further comprising:

means for retrieving a stored unit of data specified by the storage request;

and

means for verifying the correctness of the stored unit of data; and

upon detection of an error in the retrieved unit of data, means for retrieving the correct unit of data using data stored in the others of the arbitrary subset of the multiple storage nodes.

39. (new) A system for managing data storage in a network comprising multiple storage nodes, the method comprising:

means for communicating a storage request to at least one storage node;

means for causing the at least one storage node to implement the storage request using an arbitrary subset of the storage nodes;

means for attempting to retrieve the stored unit of data from the arbitrary subset of the multiple storage nodes;

Serial No. 09/777,139

Reply to Final Office Action of January 5, 2006

means for detecting unavailability of one or more network storage nodes;
and

in response to detected unavailability, means for retrieving the correct unit of data using data stored in others of the arbitrary subset of the multiple storage nodes.

40. (new) The system of claim 39, wherein the unavailability is caused by failure of one or more of the network storage nodes.

41. (new) The system of claim 39, wherein the unavailability is caused by congestion/failure of a network link leading to one or more of the network storage nodes.